arXiv:physics/0605065v1 [physics.soc-ph] 7 May 2006

Predictability, Risk and Online Management in a Complex System of Adaptive Agents

David M.D. Smith and Neil F. Johnson
Physics Department, Oxford University, Oxford OX1 3PU, U.K.

February 2, 2008

Abstract

We discuss the feasibility of predicting, managing and subsequently manipulating, the future evolution of a Complex Adaptive System. Our archetypal system mimics a population of adaptive, interacting objects, such as those arising in the domains of human health and biology (e.g. cells), financial markets (e.g. traders), and mechanical systems (e.g. robots). We show that short-term prediction yields corridors along which the model system will, with high probability, evolve. We show how the widths and average direction of these corridors varies in time as the system passes through regions, or *pockets*, of enhanced predictability and/or risk. We then show how small amounts of 'population engineering' can be undertaken in order to steer the system away from any undesired regimes which have been predicted. Despite the system's many degrees of freedom and inherent stochasticity, this dynamical, 'soft' control over future risk requires only minimal knowledge about the underlying composition of the constituent multi-agent population.

1 Introduction

Complex Adaptive Systems (CAS) are of great theoretical interest because they comprise large numbers of interacting objects or 'agents' which, unlike particles in traditional physics, change their behaviour based on experience [1]. Such adaptation yields complicated feedback processes at the microscopic level, which in turn generate complicated global dynamics at the macroscopic level. CAS also arguably represent the 'hard' problem in biology, engineering, computation and sociology[1]. Depending on the application domain, the agents in CAS may be taken as representing species, people, cells, computer hardware or software, and are typically quite numerous, e.g. $10^2 - 10^3$[1, 2].

There is also great practical interest in the problem of predicting and subsequently controlling a Complex Adaptive System. Consider the enormous

task facing a Complex Adaptive Systems 'manager' in charge of overseeing some complicated computational, biological, medical, sociological or even economic system. He would certainly like to be able to predict its future evolution with sufficient accuracy that he could foresee the system heading towards any 'dangerous' areas. However, prediction is not enough – he also needs to be able to steer the system away from this dangerous regime. Furthermore, the CAS manager needs to be able to achieve this *without* detailed knowledge of the present state of its thousand different components, nor does he want to have to shut down the system completely. Instead he is seeking for some form of 'soft' control which can be applied 'online' while the system is still evolving.

Such online management of a Complex System therefore represents a significant theoretical and practical challenge. However, the motivation for pursuing such a goal is equally great given the wide range of important real-world systems that can be regarded as complex – from engineering systems through to human health, social systems and even financial systems.

In purely deterministic systems with only a few degrees of freedom, it is well known that highly complex dynamics such as chaos can arise [3] making any control very difficult. The 'butterfly effect' whereby small perturbations can have huge uncontrollable consequences, comes to mind. One would think that things would be considerably worse in a CAS, given the much larger number of interacting objects. As an additional complication, a CAS may also contain stochastic processes at the microscopic and/or macroscopic levels, thereby adding an inherently random element to the system's dynamical evolution. The Central Limit Theorem tells us that the combined effect of a large number of stochastic processes tends fairly rapidly to a Gaussian distribution. Hence, one would think that even with reasonably complete knowledge of the present and past states of the system, the evolution would be essentially diffusive and hence difficult to control without imposing substantial global constraints.

In this paper, we address this question of dynamical control for a simplified yet highly non-trivial model of a CAS. We show that a surprising level of prediction and subsequent control can be achieved by introducing small perturbations to the agent heterogeneity, i.e. 'population engineering'. In particular, the system's global evolution can be managed and undesired future scenarios avoided. Despite the many degrees of freedom and inherent stochasticity both at the microscopic and macroscopic levels, this global control requires only minimal knowledge on the part of the 'system manager'. For the somewhat simpler case of Cellular Automata, Israeli and Goldenfeld[4] have recently obtained the remarkable result that computationally irreducible physical processes can become computationally reducible at a coarse-grained level of description. Based on our findings, we speculate that similar ideas may hold for a far wider class of system comprising populations of decision-taking, adaptive agents.

It is widely believed (see for example, Ref. [5]) that Arthur's so-called El Farol Bar Problem [6, 7] provides a representative toy model of a CAS where objects, components or individuals compete for some limited global resource (e.g. space in an overcrowded area). To make this model more complete in terms of real-world complex systems, the effect of network interconnections has recently been incorporated [8, 9, 10, 11]. The El Farol Bar Problem concerns the collective decision-making of a group of potential bar-goers (i.e. agents) who use limited global information to predict whether they should attend a potentially overcrowded bar on a given night each week. The Statistical Mechanics community has adopted a binary version of this problem, the so-called Minority Game (MG) (see Refs.[12] through [25]), as a new form of Ising model which is worthy of study in its own right because of its highly non-trivial dynamics. Here we consider a general version of such multi-agent games which (a) incorporates a finite time-horizon H over which agents remember their strategies' past successes, to reflect the fact that the more recent past should have more influence than the distant past, and (b) allows for fluctuations in agent numbers, since agents only participate if they possess a strategy with a sufficiently high success rate[15]. The formalism we employ is applicable to any CAS which can be mapped onto a population of N objects repeatedly taking actions in the form of some global 'game'.

The paper has several parts. Initially (section 2), we discuss a very simple two state 'game' to introduce and familiarise the reader to the nature of the mathematics which is explored in further detail in the rest of the paper. We then more formally establish a common framework for describing the spectrum of future paths of the complex adaptive system (section 3). This framework is general to any complex system which can be mapped onto a general B-A-R (Binary Agent Resource) model in which the system's future evolution is governed by past history over an arbitrary but finite time window H (the $Time - Horizon$). In fact, this formalism can be applied to any CAS whose internal dynamics are governed by a Markov Process[15], providing the tools whereby we can monitor the future evolution both with and without the perturbations to the population's composition. In section 4, we discuss the B-A-R model in more detail, further information is provided in [26]. We emphasize that such B-A-R systems are *not* limited to the well-known El Farol Bar Problem and Minority Games - instead these two examples are specific limiting cases. Initial investigations of a finite time-horizon version of the Minority Game were first presented in [15]. In section 5, we consider the system's evolution in the absence of any such perturbations, hence representing the system's natural evolution. In section 6, we revisit this evolution in the presence of control, where this control is limited to relatively minor perturbations at the level of the heterogeneity of the population. In section 7 we revisit the toy model of section 2 to provide a reduced form of formalism for generating averaged quantities of the future possibilities. In section 8 we discuss concluding remarks and possible

extensions.

2 A Tale of Two Dice

In this section we examine a very simple toy model employed to generate a time series (analogue output) and introduce the Future-Cast formalism to describe the model's properties. This toy model comprises two internal states, $\boldsymbol{A}$ and $\boldsymbol{B}$, and two dice also denoted $\boldsymbol{A}$ and $\boldsymbol{B}$. We make these dice generic in that we assign their faces values and these are not equal in likelihood. The rules of the model are very simple. When the system is in state $\boldsymbol{A}$, dice $\boldsymbol{A}$ is rolled and similarly for dice $\boldsymbol{B}$. The outcome, δ_t, of the relevant dice is used to increment a time (price) series, whose update can be written

$$S_{t+1} = S_t + \delta_t \tag{1}$$

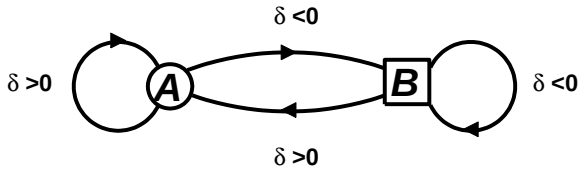

Figure 1: The internal state transitions on the De Bruijn graph.

The model employs a very simple rule to govern the transitions between its internal states. If the outcome δ_t is greater than zero (recall that we have re-assigned the values on the faces) the internal state at time $t+1$ is $\boldsymbol{A}$ and consequently dice $\boldsymbol{A}$ will be used at the next step regardless of the dice used at this time step. Conversely, if $\delta_t < 0$, the internal state at $t+1$ will be $\boldsymbol{B}$ and dice $\boldsymbol{B}$ used for the next increment[1]. These transitions are shown in figure 1.

Let us prescribe our dice some values and observe the output of the system, namely rolling dice $\boldsymbol{A}$ could yield values $\{-5, +3, +8\}$ with probabilities $\{0.3, 0.5, 0.2\}$ and $\boldsymbol{B}$ yields values $\{-10, -8, +3\}$ with probabilities $\{0.1, 0.5, 0.4\}$. These are shown in figure 2. Let us consider that at some time t the system is in state $\boldsymbol{A}$ with the value of the time series being $S(t)$ and we wish to investigate the possible output over the next U time-steps ($S(t+U)$). Some examples of the system's change in output over the next 10 time-steps are shown in figure 3. The circles represent the system being in state $\boldsymbol{A}$ and the squares the system in $\boldsymbol{B}$.

[1]This state transition rule is not critical to the formalism. For example, the rule could be to use dice $\boldsymbol{A}$ if the last increment was odd and $\boldsymbol{B}$ is even, or indeed any property of the outcomes.

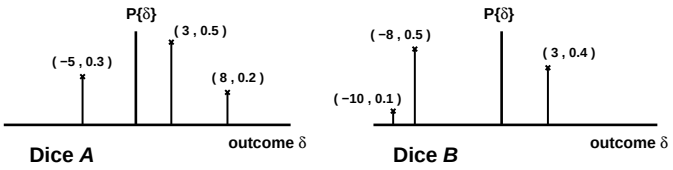

Figure 2: The outcomes and associated probabilities of our two dice.

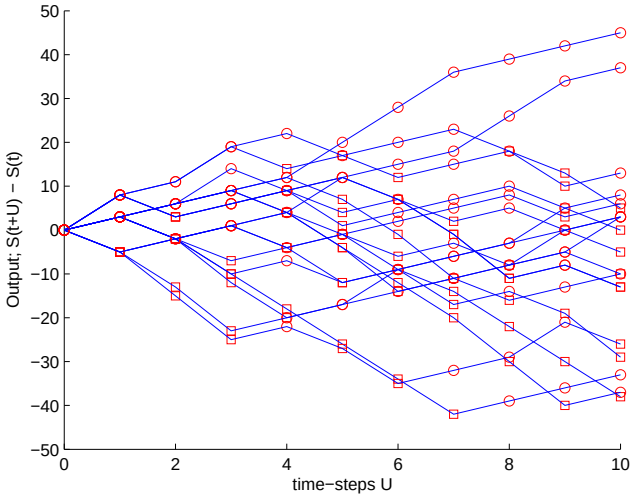

Figure 3: The toy model's output over 10 time-steps for 20 realisations where the state at time t is $\boldsymbol{A}$ for all paths.. The circles represent the system being in state $\boldsymbol{A}$ and the squares the system in $\boldsymbol{B}$.

The stochasticity inherent in the system is evident in figure 3. Many possible paths could be realised even though they all originate from state $\boldsymbol{A}$, and only a few of them are depicted. If one wanted to know more about system's output after these 10 time-steps, one could look at many more runs and look at a histogram of the output for $U = 10$. This Monte-Carlo[28] technique has been carried out in figure 4. It denotes the possible change in value of the systems analogue output after 10 time-steps and associated probability derived from 10^6 runs of the system all with the same initial starting state ($\boldsymbol{A}$) at time t. However, this method is a numerical approximation. For accurate analysis of the system over longer time scales, or for more complex systems, it might prove both inaccurate and/or computationally intensive.

For a more detailed analysis of our model, we must look at the internal state dynamics and how they map to the output. Let us consider all

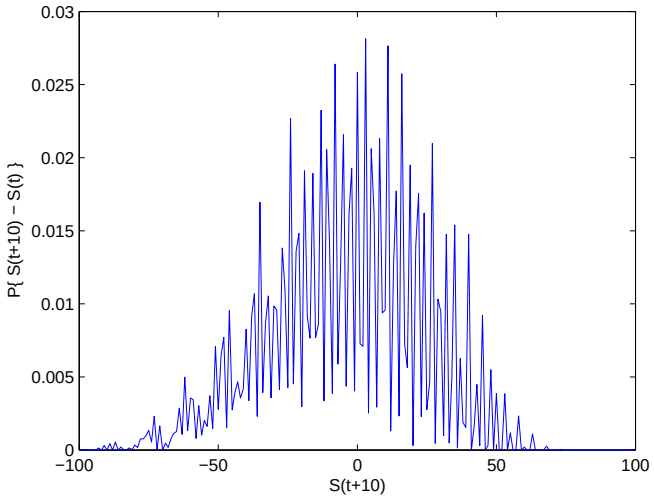

Figure 4: The change in the system's output $S(t+U) - S(t)$ at $U = 10$, and associated probability as calculated for 10^6 time-series realisations.

possible eventualities over 2 time-steps, again starting in state $\boldsymbol{A}$ at time t. All possible paths are denoted on figure 5. The resulting possible values of change in the output at time $t+2$ time-steps and their associated probabilities are given explicitly too. These values are exact in that they are calculated from the dice themselves. The $Future - Cast$ frame work which we now introduce will allow us to perform similar analysis over much longer periods. First consider the possible values of $S(t+2) - S(t)$, which result in the system being in state $\boldsymbol{A}$. The state transitions that could have occurred for these to arise are $\boldsymbol{A} \rightarrow \boldsymbol{A} \rightarrow \boldsymbol{A}$ or $\boldsymbol{A} \rightarrow \boldsymbol{B} \rightarrow \boldsymbol{A}$. The paths following the former can be considered a convolution (explores all possible paths and probabilities[2]) of the distribution of possible values of $S(t+1) - S(t)$ in state $\boldsymbol{A}$ with the distribution corresponding to the $\boldsymbol{A} \rightarrow \boldsymbol{A}$ transition. Likewise, the latter a convolution of the distribution of $S(t+1) - S(t)$ in state $\boldsymbol{B}$ with the distribution corresponding to a $\boldsymbol{B} \rightarrow \boldsymbol{A}$ transition. The resultant distribution of possibilities in state $\boldsymbol{A}$ at time $t+2$ is just the superposition of these convolutions as described in figure 6. We can set the initial value $S(t)$ to zero because the absolute position has no bearing on the evolution on the system. As such $S(t+U) - S(t) \equiv S(t+U)$.

So we note that if at some time $t+U$ there exists a distribution of values in our output which are in state $\boldsymbol{A}$ then the convolution of this with the $\boldsymbol{B} \rightarrow \boldsymbol{A}$ transition distribution will result in a contribution at $t+U+1$ to our

[2]We define and use the discrete convolution operator $\otimes$ such that $(f \otimes g)\,|_i = \sum_{j=-\infty}^{\infty} f(i-j)g(i)$.

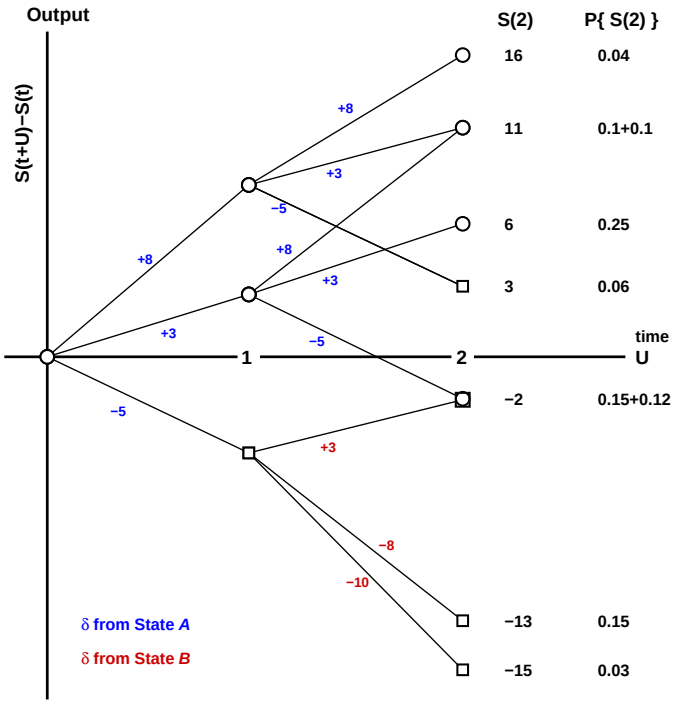

Figure 5: All possible paths after 2 time-steps and associated probability.

distribution of S which will also be in $\boldsymbol{A}$. This is evident in the right hand side of 6. Let us define all these distributions. We denote all possible values in our output U time-steps beyond the present (time t) that are in state $\boldsymbol{A}$ as the function $\varsigma_A^U(S)$ and $\varsigma_B^U(S)$ is the function that describes the values of our output time series that are in state $\boldsymbol{B}$ (as shown in figure 6). We can similarly describe the transitions as prescribed by our dice. The possible values allowed and corresponding likelihoods for the $\boldsymbol{A} \to \boldsymbol{A}$ transition are denoted $\Upsilon_{A\to A}(\delta)$ and similarly $\Upsilon_{A\to B}(\delta)$ for $\boldsymbol{A} \to \boldsymbol{B}$, $\Upsilon_{B\to A}(\delta)$ for $\boldsymbol{B} \to \boldsymbol{A}$ and $\Upsilon_{B\to B}(\delta)$ for the $\boldsymbol{B} \to \boldsymbol{B}$ state change. These are shown in figure 7.

We can now construct the Future-Cast. We can express the evolution of the output in their corresponding states as the superposition of the required convolutions:

$$\begin{aligned} \varsigma_A^{U+1}(S) &= \Upsilon_{A\to A}(\delta) \otimes \varsigma_A^U(S) + \Upsilon_{B\to A}(\delta) \otimes \varsigma_B^U(S) \\ \varsigma_B^{U+1}(S) &= \Upsilon_{A\to B}(\delta) \otimes \varsigma_B^U(S) + \Upsilon_{B\to B}(\delta) \otimes \varsigma_B^U(S) \end{aligned} \qquad (2)$$

where $\otimes$ is the discrete convolution operator as defined in footnote 2. Recall that $S(0)$ has been set to zero such that we can consider the possible changes

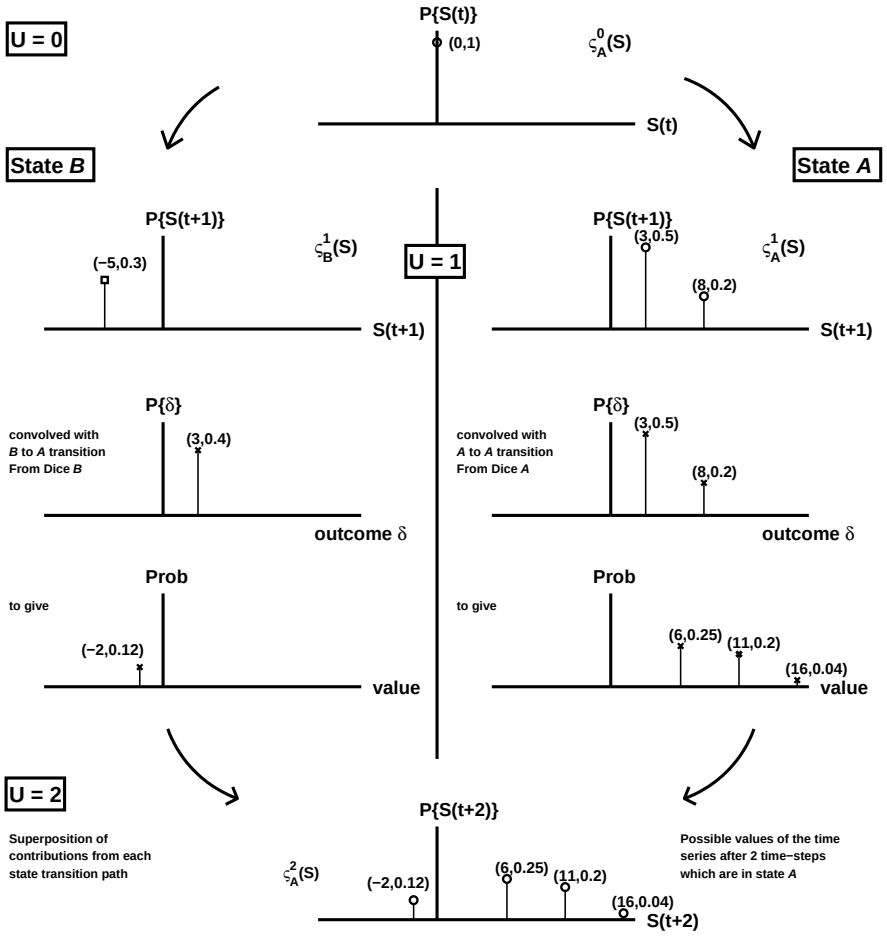

Figure 6: All possible paths after 2 time-steps which could result in the system being in state $\boldsymbol{A}$.

in output and the output itself to be identical. We can write this more concisely:

$$\underline{\varsigma^{U+1}} \;=\; \underline{\underline{\Upsilon}}\;\underline{\varsigma^{U}} \tag{3}$$

Where the element $\Upsilon_{1,1}$ contains the function $\Upsilon_{A\to A}(\delta)$ and the operator $\otimes$ and the element ς_1^{U+1} is the distribution $\varsigma_A^{U+1}(S)$. We note that this matrix of functions and operators is static, so only need computing once. As such we can rewrite equation 3 as

$$\underline{\varsigma^{U}} \;=\; \underline{\underline{\Upsilon}}^{U}\;\underline{\varsigma^{0}} \tag{4}$$

such that $\underline{\varsigma^0}$ contains the state-wise information of our starting point (time t). This is the Future-Cast process. For a system starting in state $\boldsymbol{A}$ and

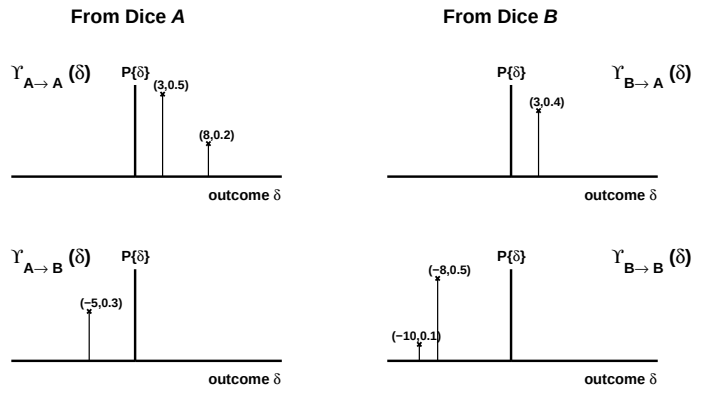

Figure 7: The state transition distributions as prescribed by our dice.

with the start value of the time series of zero, the elements of $\underline{\varsigma^0}$ are as shown in figure 8. Applying the Future-Cast process, we can look precisely at the

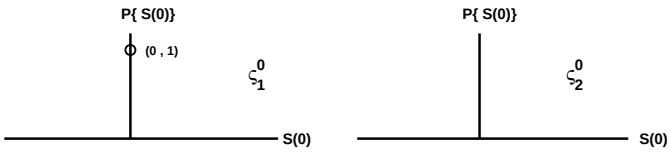

Figure 8: The initial elements of our output distributions vector, $\underline{\varsigma^0}$ when starting the system in state $\boldsymbol{A}$ with initial value of zero in our time series.

systems potential output at any number of time-steps into the future. If we wish to consider the process over 10 steps again, we applying as following:

$$\underline{\varsigma^{10}} \;=\; \underline{\underline{\Upsilon}}^{10}\;\underline{\varsigma^0} \tag{5}$$

The resultant distribution of possible outputs (we denote $\Pi^U(S)$) is then just the superposition of the contributions from each state. This distribution of the possible outputs at some time into the future is what we call the $Future-Cast$.

$$\Pi^{10} \;=\; \varsigma_1^{10} \;+\; \varsigma_2^{10} \tag{6}$$

This leads to distribution shown in figure 9. Although the exact output as calculated using the Future-Cast process demonstrated in figure 9 compares well with the brute force numerical results of the Monte-Carlo technique in figure 4, it allows us to perform some more interesting analysis without

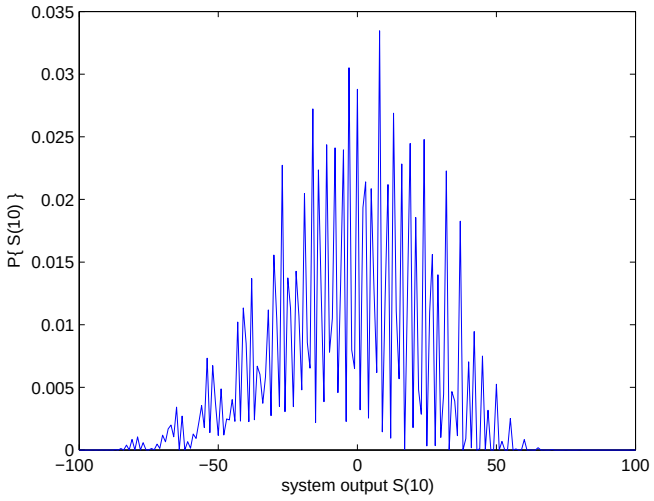

Figure 9: The actual probability distribution $\Pi^{10}(S))$ of the output after 10 time-steps starting the system in state $\boldsymbol{A}$ with initial value of zero in our time series.

much more work, let alone computational exhaustion. Consider that we don't know the initial state of the system or that we want to know characteristic properties of the system. This might include wanting to know what the system does, on average over one time-step increments. We could for example look run the system for a very long time and investigate a histogram of the output movements over single time-steps as shown in figure 10. However, we can use the formalism to generate this result exactly. Imagine that model has been run for a long time but we don't know which state it is in. Using the probabilities associated with the transitions between the states, we can infer the likelihood that the system's internal state is either $\boldsymbol{A}$ or $\boldsymbol{B}$. Let the element Γ_1^t represent the probability that the system is in state $\boldsymbol{A}$ at some time t and Γ_2^t that the system is in state $\boldsymbol{B}$. We can express these quantities at time $t+1$ by considering the probabilities of going between the states:

$$\begin{aligned} \Gamma_1^{t+1} &= T_{A\to A}\Gamma_1^t + T_{B\to A}\Gamma_2^t \\ \Gamma_2^{t+1} &= T_{A\to B}\Gamma_2^t + T_{B\to B}\Gamma_2^t \end{aligned} \tag{7}$$

Where $T_{A\to A}$ represents the probability that when the system is in state $\boldsymbol{A}$ it will be in state $\boldsymbol{A}$ at the next time-step. We can express this more concisely:

$$\underline{\Gamma^{t+1}} = \underline{\underline{T}}\,\underline{\Gamma^t} \tag{8}$$

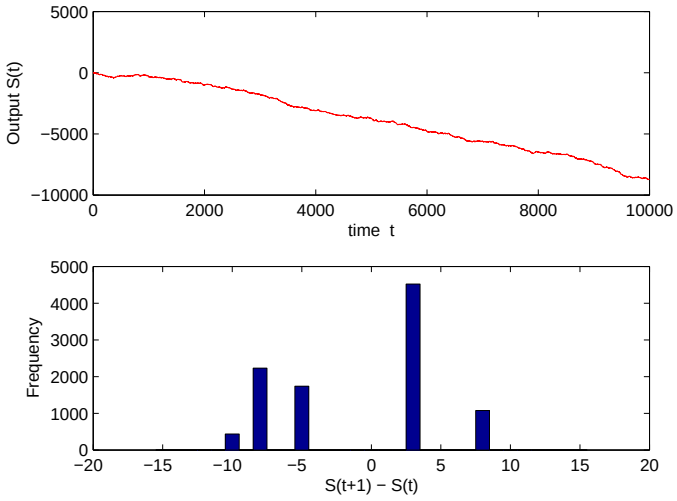

Figure 10: The one step increments of our time-series as run over 10000 steps.

Such that $T_{1,1}$ is equivalent to $T_{A\to A}$. This is a Markov Chain. From the nature of our dice, we can trivially calculate the elements of $\underline{\underline{T}}$. The value for $T_{A\to A}$ is the sum over all elements in $\Upsilon_{A\to A}(\delta)$ or more precisely:

$$T_{A\to A} = \sum_{\delta=-\infty}^{\infty} \Upsilon_{A\to A}(\delta) \tag{9}$$

The Markov Chain transition matrix $\underline{\underline{T}}$ for our system can thus be trivially written:

$$\underline{\underline{T}} = \begin{pmatrix} 0.7 & 0.4 \\ 0.3 & 0.6 \end{pmatrix} \tag{10}$$

The static probabilities of the system being in either of its two states are given by the eigenvector solution to equation 11 with eigenvalue 1. This is equivalent to looking at the relative occurrence of the two states if the system were to be run over infinite time.

$$\underline{\Gamma} \;=\; \underline{\underline{T}}\,\underline{\Gamma} \tag{11}$$

For our system, the static probability associated with being in state $\boldsymbol{A}$ is $\frac{4}{7}$ and obviously $\frac{3}{7}$ for state $\boldsymbol{B}$. To look at the characteristic properties we are going to construct an initial vector similar to ς^0 in equation 2) for our Future-Cast formalism to act on but which is related to the static probabilities contained by the solution to equation 11. This vector is denoted $\underline{\kappa}$ and its

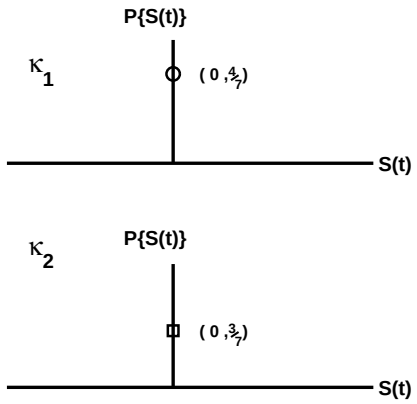

Figure 11: Explicit depiction of the elements of vector $\underline{\kappa}$ which is used to analyse characteristic behaviour of the system.

form is described in figure 11. We use employ $\underline{\kappa}$ in the Future-Cast over one time-step as in equation 12:

$$\underline{\varsigma^1} \;=\; \underline{\underline{\Upsilon}}\,\underline{\kappa} \tag{12}$$

The resulting distribution of possible outputs from superimposing the elements of $\underline{\varsigma}^1$ is the exact representation of the one time-step increments $(S(t+1) \; - S(t))$ of the system if it were allowed to be run infinitely. We call this characteristic distribution Π^1_{char}. Applying the process a number of times will yield the exact distributions Π^U_{char} equivalent to looking at all values of $S(t+U) \; - S(t)$ which is a rolling window length U over an infinite time series as in figure 12. This is also equivalent to running the system forward in time U time-steps from unknown initial state, investigating all possible paths. The Markovian nature of the system means that this is not the same as the convolution of the one time-step characteristic Future-Cast convolved with it self U times.

Clearly the characteristic Future-Cast over one time-step in figure 12 compares well with that of figure 10.

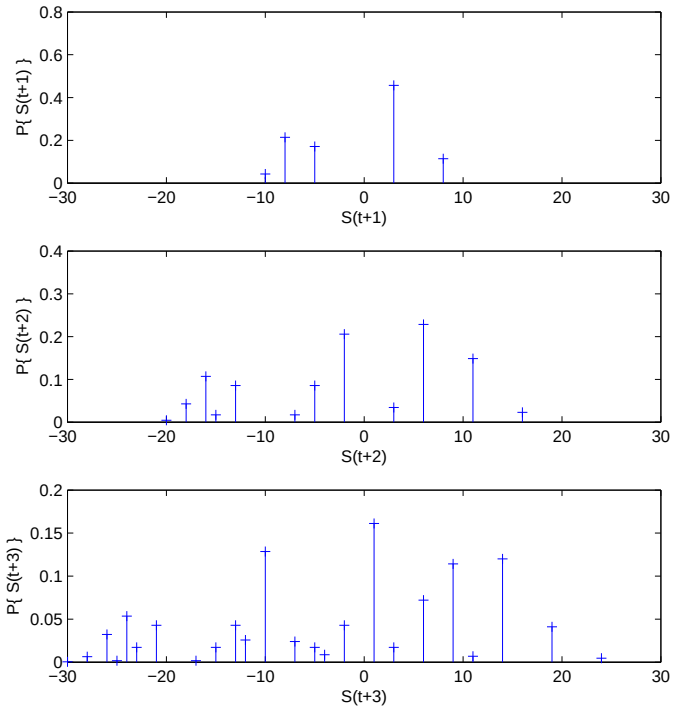

Figure 12: The characteristic behaviour of the system for $U = 1, 2, 3$ time-steps into the future from unknown initial state. This is equivalent to looking at the relative frequency of occurrence of changes in output values over 1,2, and 3 time-step rolling windows

3 The Evolution of the Complex Adaptive System

Here we provide a general formalism applicable to any Complex System which can be mapped onto a population of N species or 'agents' who are repeatedly taking actions in some form of global 'game'. At each time-step each agent makes a (binary) decision $a_{\mu(t)}$ in response to the global information $\mu(t)$ which may reflect the history of past global outcomes. This global information is of the form of a bitstring of length m. For a general game, there exists some winning outcome $w(t)$ based on the aggregate action of the agents. Each agent holds a subset of all possible strategies - by assigning this subset randomly to each agent, we can mimic the effect of large-scale heterogeneity in the population. In other words, we have a simple way of generating a potentially diverse ecology of species, some of which may be similar but others quite different. One can hence investigate a typically-

diverse ecology whereby all possible species are represented, as opposed to special cases of ecologies which may themselves generate pathological behaviour due to their lack of diversity.

The aggregate action of the population at each time-step t is represented by $D(t)$, which corresponds to the accumulated decisions of all the agents and hence the (analogue) output variable of the system at that time-step. The goal of the game, and hence the winning decision, could be to favour the minority group (MG), the majority group or indeed any function of the macroscopic or microscopic variables of the system. The individual agents do not themselves need to be conscious of the precise nature of the game, or even the algorithm for deciding how the winning decision is determined. Instead, they just know the global outcome, and hence whether their own strategies predicted the winning action[3]. The agents then reward the strategies in their possession if the strategy's predicted action would have been correct if that strategy was implemented. The global history is then updated according to the winning decision. It can be expressed in decimal form as follows:

$$\mu(t) = \sum_{i=1}^{m} 2^{i-1}[w(t-i)+1] \tag{13}$$

The system's dynamics are defined by the rules of the game. We will consider here the class of games whereby each agent uses his highest-scoring strategy at each timestep, and agents only participate if they possess a strategy with a sufficiently high success rate. [N.B. Both of these assumptions can be relaxed, thereby modifying the actual game being played]. The following two scenarios might then arise during the system's evolution:

- An agent has two (or more) strategies which are tied in score and are above the confidence level, and the decisions from them differ.
- The number of agents choosing each of the two actions is equal, hence the winning decision is undecided.

We will consider these cases to be resolved with a fair 'coin toss', thereby injecting stochasticity or 'noise' into the system's dynamical evolution. In the first case, each agent will toss his own coin to break the tie, while in the second the Game-master tosses a single coin. To reflect the fact that evolving systems will typically be non-stationary, and hence the more distant past will presumably be perceived as less relevant to the agents, the strategies are rewarded as to whether they would have made correct predictions over the last H time-steps of the game's running. There is no limit on the size of H other than it is finite and constant. The time-horizon represents a trajectory of length H on the de Bruijn graph in $\mu(t)$ (history) space[15] as shown in

[3]The algorithm used by the 'Game-master' to generate the winning decision could also incorporate a stochastic factor.

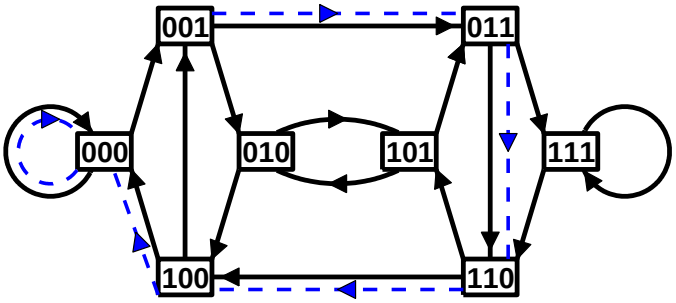

Figure 13: A path of time-horizon length $H = 5$ (dashed line) superimposed on the de Bruin graph for $m = 3$. The 8 global outcome states represent the 8 possible bitstrings for the global information, and correspond to the global outcomes for the past $m = 3$ timesteps.

figure 13. The stochasticity in the game means that for a given time-horizon H and a given strategy allocation in the population, the output of the system is not always unique. We will denote the set of all possible outputs from the game at some number of time-steps beyond the time-horizon H, as the Future-Cast.

It is useful to work in a time-horizon space $\underline{\Gamma_t}$ of dimension 2^{m+H}. An element Γ_t corresponds to the last $m+H$ elements of the bitstring of global outcomes (or equivalently, the winning actions) produced by the game. This dimension is constant in time whereas for a non-time-horizon game it would grow linearly. For any given time-horizon state, Γ_t, there exists a unique score vector $\underline{G(t)}$ which is the set of scores $G_R(t)$ for all the strategies which an agent could possess. As such, for each particular time-horizon state, there exists a unique probability distribution of the aggregate action, $D(t)$. This distribution of possible actions when a specified state is reached will necessarily be the same each time that state is revisited. Thus, it is possible to construct a transition matrix (c.f. Markov Chain[15]) $\underline{\underline{T}}$ of probabilities for the movements between these time-horizon states such that $\underline{P(\Gamma_t)}$ can be expressed as

$$\underline{P(\Gamma_t)} \;=\; \underline{\underline{T}}\;\underline{P(\Gamma_{t-1})} \tag{14}$$

where $\underline{P(\Gamma_t)}$ is a vector of dimension 2^{m+H} containing the probabilities of being in a given state Γ at time t

The transition matrix of probabilities is constant in time and necessarily sparse. For each state, there are only two possible winning decisions. The number of non-zero elements in the matrix is thus $\leq 2^{(m+H+1)}$. We can use the transition matrix in an eigenvector-eigenvalue problem to obtain the stationary state solution of $\underline{P(\Gamma)} = \underline{\underline{T}}\;\underline{P(\Gamma)}$. This also allows calculation of

some time-averaged macroscopic quantities of the game [15][4].

To generate the Future-Cast, we want to calculate the quantities in output space. To do this, we require;

- The probability distribution of $D(t)$ for a given time-horizon;
- The corresponding winning decisions, $w(t)$, for given $D(t)$;
- An algorithm generating output in terms of $D(t)$.

To implement the Future-Cast, we need to map from the transitions in the state space internal to the system to the macroscopic observables in the output space (often cumulative excess demand). We know that in the transition matrix, the probabilities represent the summation over a distribution of possible aggregate actions which is binomial in the case where the agents are limited to two possible decisions. Using the output generating algorithm, we can construct an 'adjacency' matrix $\underline{\underline{\Upsilon}}$ analogous to the transition matrix $\underline{\underline{T}}$, with the same dimensions. The elements of $\underline{\underline{\Upsilon}}$, contain probability distribution functions of change in output corresponding to the non-zero elements of the transition matrix together with the discrete convolution operator $\otimes$ whose form depends on that of the output generating algorithm.

The adjacency matrix of functions and operators can then be applied to a vector, $\underline{\varsigma^{U=0}(S)}$, containing information about the current state of the game and of the same dimension as $\underline{\Gamma_t}$. $\underline{\varsigma^{U=0}(S)}$ not only describes the time-horizon state positionally through its elements but also the current value in the output quantity S within that element. At $U = 0$, the state of the system is unique so there is only one non-zero element within $\underline{\varsigma^{U=0}(S)}$. This element corresponds to a probability distribution function of the current output value, its position within the vector corresponding to the current time-horizon state. The probability distribution function is necessarily of value unity at the current value or, for a Future-Cast expressed in terms of change in output from the current value, unity at the origin. The Future-Cast process for U time-steps beyond the present state can then be described by

$$\underline{\varsigma^{U}(S)} \;=\; \underline{\underline{\Upsilon}}^{U}\,\underline{\varsigma^{0}(S)} \tag{15}$$

The actual Future-Cast, $\Pi(S,U)$, is then computed by superimposing the elements of the output/time-horizon state vector:

$$\Pi^{U}(S) = \sum_{i=1}^{2^{(m+H)}} \varsigma_i^{U}(S). \tag{16}$$

[4]The steady state eigenvector solution is an exact expression equivalent to pre-multiplying the probability state vector $\underline{P(\Gamma_t)}$ by $\underline{\underline{T}}^{\infty}$. This effectively results in a probability state vector which is time-averaged over an infinite time-interval.

Thus the Future-Cast, $\Pi^U(S)$, is a probability distribution of the outputs possible at U time-steps in the future.

As a result of the state dependence of the Markov Chain, Π is non-Gaussian. As with the steady-state solution of the state space transition matrix, we would like to find a 'steady-state' equivalent for the output space[5] of the form

$$\Pi^1_{char}(S) = \left\langle \Pi^1(S) \right\rangle_\infty \tag{17}$$

where the one-timestep Future-Cast is time-averaged over an infinitely long period. Fortunately, we have the steady state solutions of $\underline{P(\Gamma)} = \underline{\underline{T}}\ \underline{P(\Gamma)}$ which are the (static) probabilities of being in a given time-horizon state at any time. By representing these probabilities as the appropriate functions, we can construct an 'initial' vector, $\underline{\kappa}$, similar in form to $\underline{\varsigma(S,0)}$ in equation 15 but equivalent to the eigenvector solution of the Markov Chain. We can then generate the solution of equation 17 for the *characteristic* Future-Cast, Π^1_{char}, for a given initial set of strategies. An element κ_i is again a probability distribution which is simply the point $(0\ ,\ P_i(\Gamma))$, the static probability of being in the time-horizon state denoted by the elements position, i. We can then get back to the Future-Cast

$$\Pi^1_{char}(S) \ = \ \sum_{i=1}^{2^{(m+H)}} \varsigma_i^1 \quad \text{where} \quad \underline{\varsigma^1} \ = \ \underline{\underline{\Upsilon}}\ \underline{\kappa}. \tag{18}$$

We can also generate characteristic Future-Casts for any number of time-steps, U, by pre-multiplying $\underline{\kappa}$ by $\underline{\underline{\Upsilon}}^U$

$$\Pi^U_{char}(S) \ = \ \sum_{i=1}^{2^{(m+H)}} \varsigma_i^U \quad \text{where} \quad \underline{\varsigma^U} \ = \ \underline{\underline{\Upsilon}}^U\ \underline{\kappa} \tag{19}$$

We note that Π^U_{char} is not equivalent to the convolution of Π^1_{char} with itself U times and as such is not necessarily Gaussian. The characteristic Future-Cast over U time-steps is simply the Future-Cast of length U from all the 2^{m+H} possible initial states where each contribution is given the appropriate weighting factor. This factor corresponds to the probability of being in that initial state. The characteristic Future-Cast can also be expressed as

$$\Pi^U_{char}(S) \ = \ \sum_{\Gamma=1}^{2^{(m+H)}} P(\Gamma)\ \Pi^U(S) \mid \Gamma \tag{20}$$

where $\Pi^U(S) \mid \Gamma$ is a normal Future-Cast from an initial time-horizon state Γ and $P(\Gamma)$ is the static probability of being in that state at a given time.

[5]Note that we can use this framework to generate time-averaged quantities of *any* of the macroscopic quantities of the system (e.g total number of agents playing) or volatility.

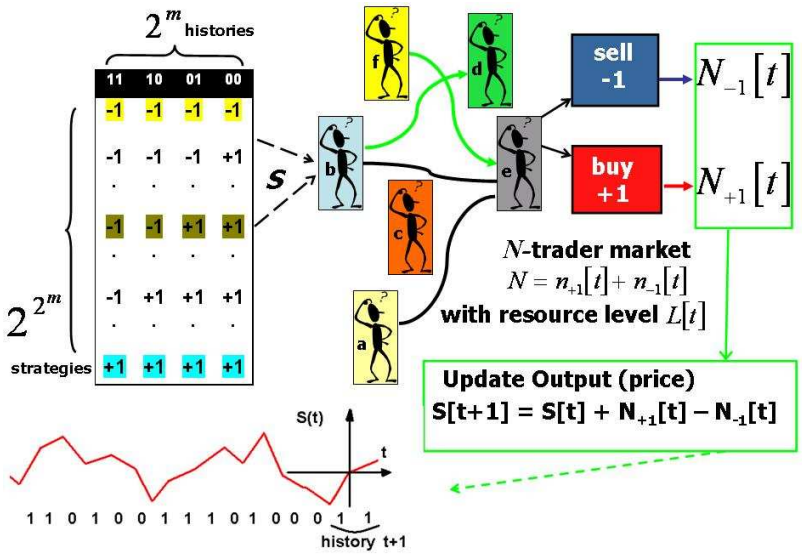

Figure 14: Schematic diagram of the Binary Agent Resource (B-A-R) system.

4 The Binary Agent Resource System

The general binary framework of the B-A-R (Binary Agent Resource) system was discussed in section 3. The global outcome of the 'game' is represented as a binary digit which favours either those choosing option $+1$ or option -1 (or equivalently 1 or 0, A or B etc.). The agents are randomly assigned s strategies at the beginning of the game. Each strategy comprises an action $a^s_{\mu(t)}$ in response to each of the 2^m possible histories μ, thereby generating a total of 2^{2^m} strategies in the Full Strategy Space [6]. At each turn of the game, the agents employ their most successful strategy, being the one with the most virtual points. The agents are thus adaptive if $s > 1$.

We have already extended the B-A-R system by introducing the time-horizon H, which determines the number of past time-steps over which virtual points are collected for each strategy. We further extend the system by the introduction of a confidence level. The agents decide whether to participate or not depending on the success of their strategies. As such, the number of active agents $N(t)$ is less than or equal to N_{tot} at any given time-step. This results in a variable number of participants per time-step

[6] We note that many features of the game can be reproduced using a Reduced Strategy Space of 2^{m+1} strategies, containing strategies which are either anti-correlated or uncorrelated with each other[12]. The framework established in the present paper is general to both the full and reduced strategy spaces, hence the full strategy space will be adopted here.

$V(t)$, and constitutes a 'Grand Canonical' game. The threshold, τ, denotes the confidence level: each agent will only participate if he has a strategy with at least r points where

$$r = T(2\tau - 1). \tag{21}$$

Agents without an active strategy become temporarily inactive.

In keeping with typical biological, ecological, social or computational systems, the Game-master takes into account a finite global resource level when deciding the winning decision at each time-step. For simplicity, we will here consider the specific case[7] whereby the resource level $L(t) = \phi V(t)$ with $0 \leq \phi \leq 1$. We denote the number of agents choosing action +1 (or equivalently A) as $N_{+1}(t)$, and those that choose action -1 (or equivalently B) as $N_{-1}(t)$. If $L(t) - N_{+1}(t) > 0$ the winning action is +1 and vice-versa. We define the winning decision 1 or 0 as follows:

$$w(t) = \text{step}[L(t) - N_{+1}(t)] \tag{22}$$

where we define step$[x]$ to be

$$\text{step}[x] = \begin{cases} 1 & \text{if } x > 0, \\ 0 & \text{if } x < 0, \\ \text{fair coin toss} & \text{if } x = 0. \end{cases} \tag{23}$$

When $x = 0$, there is no definite winning option since $N_{+1}(t) = N_{-1}(t)$, hence the Game-master uses a random coin-toss to decide between the two possible outcomes. We use a binary payoff rule for rewarding strategy scores, although more complicated versions can, of course, be used. However, we note that non-binary payoffs (e.g. a proportional payoff scheme) will decrease the probability of tied strategy scores, hence making the system more deterministic. Since we are interested in seeing the extent to which stochasticity can prevent control, we are instead interested in preserving the presence of such stochasticity. The reward function χ can be written

$$\chi[N_{+1}(t), L(t)] = \begin{cases} 1 & \text{for } w(t) = 1, \\ -1 & \text{for } w(t) = 0, \end{cases} \tag{24}$$

namely +1 for predicting the correct action and -1 for predicting the incorrect one. For a given strategy, R, the virtual points score is given by

$$G_R(t) = \sum_{i=t-T}^{t-1} a_R^{\mu(i)} \chi[N_{+1}(i), L(i)], \tag{25}$$

where $a_R^{\mu(t)}$ is the response of strategy, R, to the global information $\mu(t)$ summed over the rolling window of width H. The global output signal $D(t) = N_{+1}(t) - N_{-1}(t)$ is calculated at each iteration to generate an output time series.

[7] We note that ϕ itself could be actually be a stochastic function of the known system parameters.

5 Looking at the System's Natural Evolution

To realize all possible paths within a given game is necessarily computationally expensive. For a Future-Cast U timesteps beyond the current game state, there are necessarily 2^U winning decisions to be considered. Fortunately, not all winning decisions are realized by a specific game and the numerical generation of the Future-Cast can be made reasonably efficient.

Fortunately we can approach the Future-Cast analytically *without* having to keep track of the agents' individual microscopic properties. Instead we group the agents together via the population tensor of rank s given by $\underline{\underline{\Omega}}$, which we will refer to as the Quenched Disorder Matrix (QDM) [20]. This matrix is assumed to be constant over the time-scales of interest, and more typically is fixed at the beginning of the game. The entry $\Omega_{R2,R2,...}$ represents the number of agents holding the strategies $R1, R2, \ldots$ such that

$$\sum_{R,R',...} \underline{\underline{\Omega}}_{R,R',...} = N \tag{26}$$

For numerical analysis, it is useful to construct a symmetric version of this population tensor, $\underline{\underline{\Psi}}$. For the case $s = 2$, we will let $\underline{\underline{\Psi}} = \frac{1}{2}(\underline{\underline{\Omega}} + \underline{\underline{\Omega}}^{transpose})$ [17].

The output variable $D(t)$ can be written in terms of the decided agents $D_d(t)$ who act in a pre-determined way since they have a unique predicted action from their strategies, and the undecided agents $D_{ud}(t)$ who require an additional coin-toss in order to decide which action to take. Hence

$$D(t) = D_d(t) + D_{ud}(t). \tag{27}$$

We focus on $s = 2$ strategies per agent although the approach can be generalized. The element $\Psi_{R,R'}$ represents the number of agents holding both strategy R and R'. We can now write $D_d(t)$ as

$$D_d(t) = \sum_{R=1}^{Q} a_R^{\mu(t)} \mathcal{H}[G_R(t) - r] \sum_{R'=1}^{Q} (1 + \mathrm{sgn}[G_R(t) - G_{R'}(t)]) \Psi_{R,R'} \tag{28}$$

where Q is the size of the strategy space, $\mathcal{H}$ is the Heaviside function and $\mathrm{sgn}[x]$ is defined as

$$\mathrm{sgn}[x] = \begin{cases} 1 & \text{if } x > 0, \\ -1 & \text{if } x < 0, \\ 0 & \text{if } x = 0. \end{cases} \tag{29}$$

The volume $V(t)$ of active agents can be expressed as

$$V(t) \;=\; \sum_{R,R'} \mathcal{H}[G_R(t) - r] \big\{ sgn[G_R(t) - G_{R'}(t)] + \frac{1}{2} \delta[G_R(t) - G_{R'}(t)] \big\} \Psi_{R,R'} \tag{30}$$

where δ is the Dirac delta. The number of undecided agents $N_{ud}(t)$ is given by

$$N_{ud}(t) \;=\; \sum_{R,R'} \mathcal{H}[G_R(t)-r]\delta(G_R(t)-G_{R'}(t))[1-\delta(a_R^{\mu(t)}-a_{R'}^{\mu(t)})]\Psi_{R,R'} \quad (31)$$

We note that for $s = 2$, because each undecided agent's contribution to $D(t)$ is an integer, hence the demand of all the undecided agents $D_{ud}(t)$ can be written simply as

$$D_{ud}(t) \;\epsilon\; 2\ \mathrm{Bin}\left(N_{ud}(t), \frac{1}{2}\right) - N_{ud}(t) \quad (32)$$

where $\mathrm{Bin}(n, p)$ is a sample from a binomial distribution of n trials with probability of success p.

For any given time-horizon space-state Γ_t, the score vector $\underline{G(t)}$ (i.e., the set of scores $G_R(t)$ for all the strategies in the QDM) is unique. Whenever this state is reached, the quantity $D_d(t)$ will necessarily always be the same, as will the distribution of $D_{ud}(t)$. We can now construct the transition matrix $\underline{\underline{T}}$ giving the probabilities for the movements between these time-horizon states. The element $\underline{\underline{T}}_{\Gamma_t|\Gamma_{t-1}}$ which corresponds to the transition from state Γ_{t-1} to Γ_t, is given for the (generalisable) $s = 2$ case by

$$\begin{aligned}\underline{\underline{T}}_{\Gamma_t|\Gamma_{t-1}} = \sum_{x=0}^{N_{ud}} \Bigg\{ {}^{N_{ud}}C_x (\frac{1}{2})^{N_{ud}} \delta \Big[Sgn(D_d + 2x - N_{ud} + V(1-2\phi)) \;+\; \\ (2\mu_t \%2 - 1) \Big] \;+\; \\ {}^{N_{ud}}C_x (\frac{1}{2})^{(N_{ud}+1)} \delta \Big[Sgn(D_d + 2x - N_{ud} + V(1-2\phi)) \;+\; 0 \Big] \Bigg\} \end{aligned} \quad (33)$$

where N_{ud}, D_d implies $N_{ud} \mid \Gamma_{t-1}$ and $D_d \mid \Gamma_{t-1}$, V implies $V(t-1)$, ϕ sets the resource level as described earlier and $\mu_t\%2$ is the required winning decision to get from state Γ_{t-1} to state Γ_t. We use the transition matrix in the eigenvector-eigenvalue problem to obtain the stationary state solution of $\underline{P(\Gamma)} \;=\; \underline{\underline{T}}\ \underline{P(\Gamma)}$. The probabilities in the transition matrix represent the summation over a distribution which is binomial in the $s = 2$ case. These distributions are all calculated from the QDM which is fixed from the outset. To transfer to output-space, we require an output generating algorithm. Here we use the equation

$$S(t+1) = S(t) + D(t) \quad (34)$$

hence the output value $S(t)$ represents the cumulative value of $D(t)$, while the increment $S(t+1) - S(t)$ is simply $D(t)$. Again, we use the discrete

convolution operator $\otimes$ defined as

$$(f \otimes g) \mid_i \;\; = \sum_{j=-\infty}^{\infty} f(i-j) \times g(j). \tag{35}$$

The formalism could be extended for general output algorithms using differently defined convolution operators.

An element in the adjacency matrix for the $s = 2$ case can then be expressed as

$$\underline{\underline{\Upsilon}}_{\Gamma_t|\Gamma_{t-1}} = \Bigg\{ \sum_{x=0}^{N_{ud}} \Bigg((D_d + 2x - N_{ud}),$$
$${}^{N_{ud}}C_x (\frac{1}{2})^{N_{ud}} \delta \Big[Sgn(D_d + 2x - N_{ud} + V(1-2\phi)) + (2\mu_t \%2 - 1) \Big] +$$
$${}^{N_{ud}}C_x (\frac{1}{2})^{(N_{ud}+1)} \delta \Big[Sgn(D_d + 2x - N_{ud} + V(1-2\phi)) + 0 \Big] \Bigg) \Bigg\} \otimes \tag{36}$$

where N_{ud}, D_d again implies $N_{ud} \mid \Gamma_{t-1}$ and $D_d \mid \Gamma_{t-1}$, V implies $V(t-1)$, and $\mu_t \%2$ is the winning decision necessary to move between the required states. The Future-Cast and characteristic Future-Casts ($\Pi^U(S)$, Π^U_{char}) U time-steps into the future can then be computed for a given initial quenched disorder matrix (QDM).

We now consider an example to illustrate the implementation. In particular, we provide the explicit solution of a Future-Cast in the regime of small m and H, given the randomly chosen quenched disorder matrix

$$\underline{\underline{\Omega}} = \begin{pmatrix}
0 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 2 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\
0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 2 & 1 & 0 & 0 & 1 \\
0 & 1 & 0 & 0 & 2 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 \\
1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 1 \\
1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 2 & 0 \\
0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 2 & 1 & 0 & 0 & 1 \\
0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\
2 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 3 & 1 & 1 & 0 & 1 & 0 & 1 \\
0 & 2 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 3 & 2 & 0 & 0 & 0 & 1 & 1 \\
0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 2 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\
0 & 2 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1
\end{pmatrix}. \tag{37}$$

We consider the full strategy space and the the following game parameters:

Number of agents	N_{tot}	101
Memory size	m	2
Strategies per agent	s	2
Resource level	ϕ	0.5
Time horizon	H	2
Threshold	τ	0.51

The dimension of the transition matrix is thus $2^{H+m} = 16$.

$$\underline{\underline{T}} = \begin{pmatrix}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0.5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0625 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0.5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9375 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.1875 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.8125 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0.1875 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.1094 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0.8125 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.8906 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0312 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9688 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0.75 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.5 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.5 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{pmatrix}. \tag{38}$$

Each non-zero element in the transition matrix corresponds to a probability function in the output space in the Future-Cast operator matrix. Consider that the initial state is $\Gamma = 10$ i.e. the last 4 bits are $\{1001\}$ (obtained from running the game prior to the Future-Casting process). The initial probability state vector is the point (0,1) in the element of the vector $\underline{\varsigma}$ corresponding to time-horizon state 10. We can then generate the Future-Cast for given U (shown in figure 15).

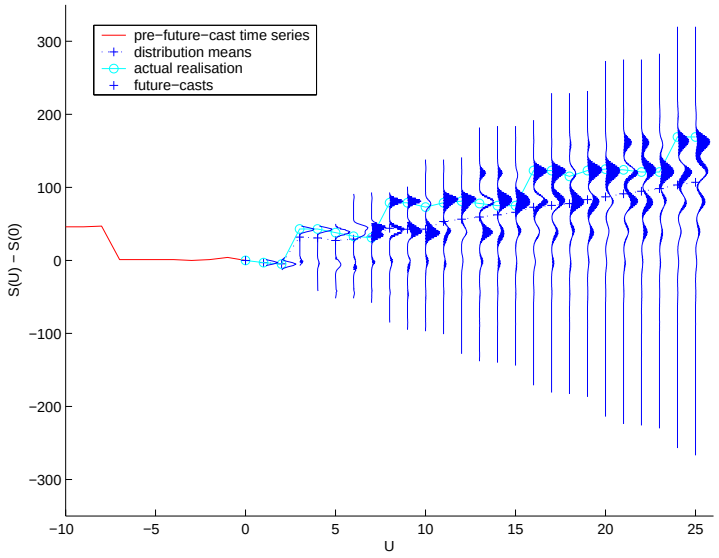

Figure 15: The (un-normalized) evolution of a Future-Cast for the given $\underline{\underline{\Omega}}$, game parameters and initial state Γ. The figure shows the last 10 time-steps prior to the Future-Cast, the means of the distributions within the Future-Cast itself, and also an actual realization of the game run forward in time.

Clearly the probability function in output space becomes smoother as U becomes larger and the number of successive convolutions increases, as highlighted by the probability distribution functions at $U = 15$ and $U = 25$

(figure 16).

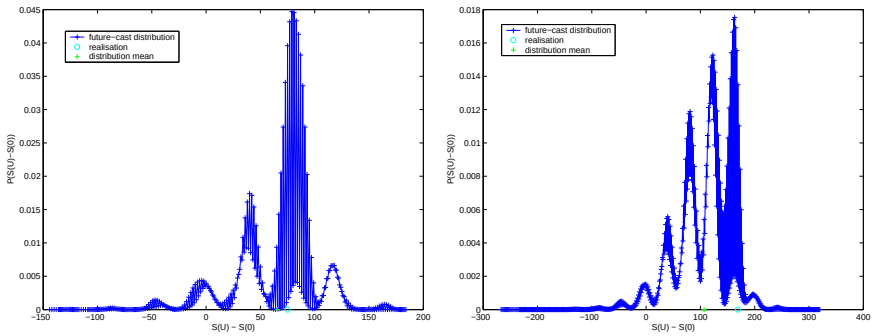

Figure 16: The probability distribution function at $U = 15, 25$ time-steps beyond the present state.

We note the non-Gaussian form of the probability distribution for the Future-Casts, emphasising the fact that such a Future-Cast approach is essential for understanding the system's evolution. An assumption of rapid diffusion toward a Gaussian distribution, and hence the future spread in paths increasing as the square-root of time, would clearly be unreliable.

6 Online Evolution Management via 'Soft' Control

For less simple parameters, the matrix dimension required for the Future-Cast process become very large very quickly. To generate a Future-Cast appropriate to larger parameters e.g. $m = 3$, $H = 10$, it is still however possible to carry out the calculations numerically quite easily. As an example, we generate a random $\underline{\underline{\Omega}}$ (the form of which is given in figure 19) and initial time-horizon appropriate to these parameters. This time-horizon is obtained by allowing the system to run prior to the Future-Cast. For visual representation reasons, the Reduced Strategy Space[20] is employed. The other game parameters are as previously stated. The game is then instructed to run down every possible winning decision path exhaustively. The spread of output at each step along each path is then convolved with the next spread such that a Future-Cast is built up along each path. Fortunately, not all paths are realized at every time-step since the stochasticity in the winning-decision/state-space results from the condition $N_{ud} \geq D_d$. The Future-Cast as a function of U and S, can thus be built up for a randomly chosen initial quenched disorder matrix (QDM) (17).

We now wish to consider the situation where it is required that the system should not behave in a certain manner. For example, it may be desirable that it avoid entering a certain regime characterised by a given

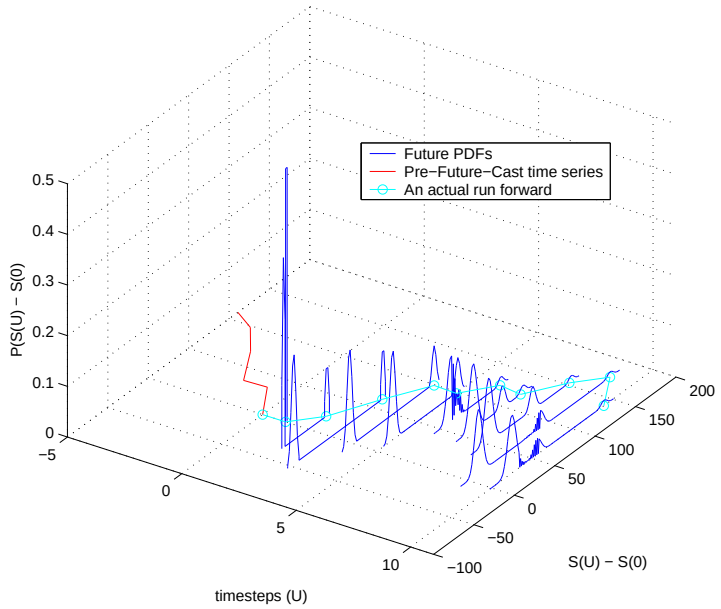

Figure 17: Evolution of $\Pi^U(S)$ for a typical quenched disorder matrix $\underline{\underline{\Omega}}$.

value of $S(t)$. Specifically, we consider the case where there is a barrier in the output space that the game should avoid, as shown in figure 18.

The evolution of the spread (i.e. standard deviation) of the distributions in time, confirms the non-Gaussian nature of the system's evolution – we note that this spread can even decrease with time[8]. In the knowledge that this barrier will be breached by this system, we therefore perturb the quenched disorder at $U = 0$. This perturbation corresponds in physical terms to an adjustment of the composition of the agent population. This could be achieved by 're-wiring' or 'reprogramming' individual agents in a situation in which the agents were accessible objects, or introducing some form of communication channel, or even a more 'evolutionary' approach whereby a small subset of species are removed from the population and a new subset added in to replace them. Interestingly we note that this 'evolutionary' mechanism need neither be completely deterministic (i.e. knowing exactly how the form of the QDM changes) nor completely random (i.e. a random perturbation to the QDM). In this sense, it seems tantalisingly close to some modern ideas of biological evolution, whereby there is some purpose mixed with some randomness.

[8]This feature can be understood by appreciating the multi-peaked nature of the distributions in question. The peaks correspond to differing paths travelled in the Future-Cast, the final distribution being a superposition of these. If these individual path distributions mean-revert, the spread of the actual Future-Cast can decrease over short time-scales.

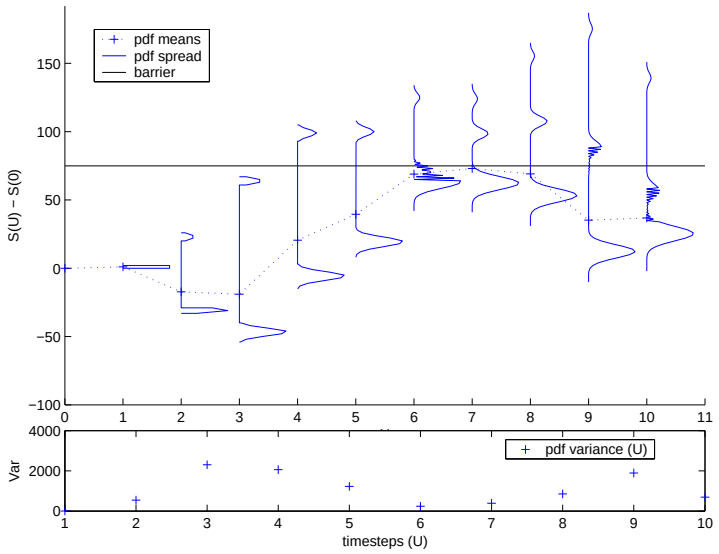

Figure 18: The evolution of the Future-Casts, and the barrier to be avoided. For simplicity the barrier is chosen to correspond to a fixed $S(t)$ value of 110, although there is no reason that it couldn't be made time-dependent. Superimposed on the (un-normalised) distributions, are the means of the Future-Casts, while their variances are shown below.

Figure 20 shows the impact of this relatively minor microscopic perturbation on the Future-Cast and global output of the system. In particular, the system has been steered away from the potentially harmful barrier into 'safer' territory.

This set of outputs is specific to the initial state of the system. More typically, we may not know this initial state. Fortunately, we can make use of the characteristic Future-Casts to make some kind of quantitative assessment of the robustness of the quenched disorder perturbation in avoiding the barrier, since this procedure provides a picture of the range of possible future scenarios.

This evolution of the characteristic Future-Casts, for both the initial and perturbed quenched disorder matrices, is shown in figure 21 A quantitative evaluation of the robustness of this barrier avoidance could then be calculated using traditional techniques of risk analysis, based on knowledge of the distribution functions and/or their low-order moments.

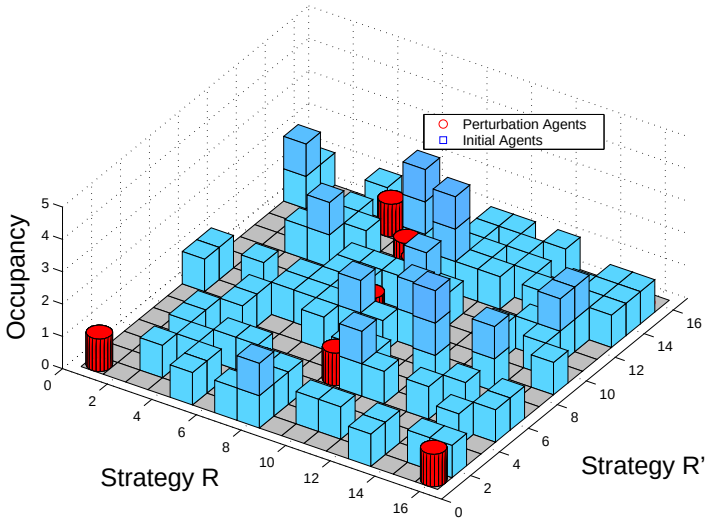

Figure 19: The initial and resulting quenched disorder matrices (QDM), shown in schematic form. The x-y axes are the strategy labels for the two strategies. The absence of a symbol denotes an empty bin (i.e. no agent holding that particular pair of strategies).

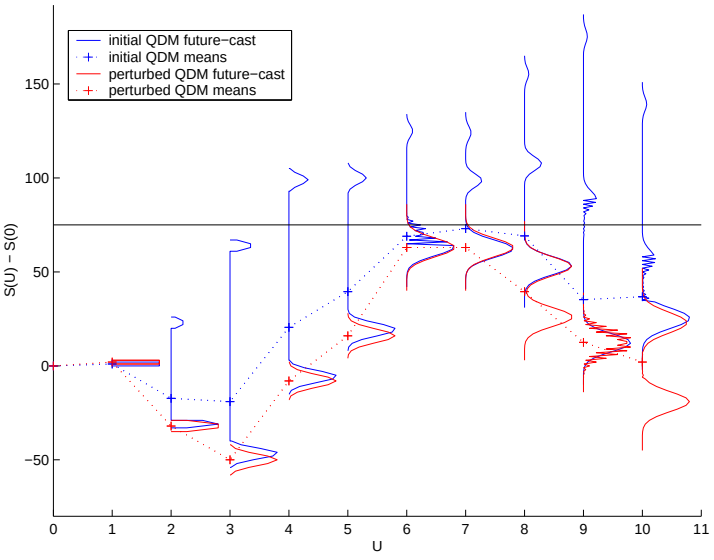

Figure 20: The evolution as a result of the microscopic perturbation to the population's composition (i.e. the QDM).

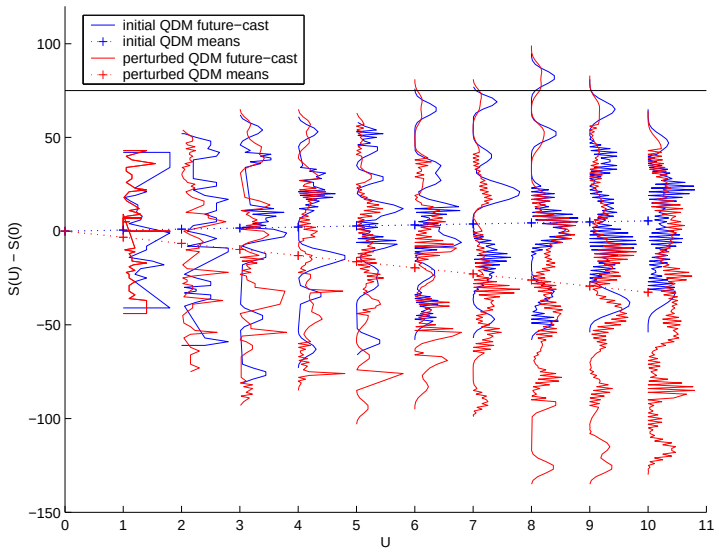

Figure 21: The characteristic evolution of the initial and perturbed QDMs.

7 A Simplified Implementation of the Future-Cast Formalism

We introduced the Future-Cast formalism to map from the internal state space of a complex system to the observable output space. Although the formalism exactly generates the probability distributions of the subsequent output from the system, it's implementation is far from trivial. This involves keeping track of numerous distributions and performing appropriate convolutions between them. Often, however it is only the lowest order moments which are of immediate concern to the system designer. Here, we show how this information can be generated without the computational exhaustion previously required. We demonstrate this procedure for a very simple two state system, although the formalism is general to a system of any number of states, governed by a Markov chain.

Recall the toy model comprising two dice of section 2 . We previously broke down the possible outputs of each according to the state transition as shown in figure 22. These distributions were used to construct the matrix $\underline{\underline{\Upsilon}}$ to form the Future-Cast process as denoted in equation 3. This acted on vector $\underline{\varsigma}^{U}$ to generate $\underline{\varsigma}^{U+1}$. The elements of these vectors contain the partial distribution of outputs which are in the state denoted by the element number at that particular time, so for the two dice model, $\varsigma_1^U(S)$ contains the distribution of output values at time $t+U$ (or U time-steps beyond the present) which correspond to the system being in state $\boldsymbol{A}$ and $\varsigma_2^U(S)$ contains those for state $\boldsymbol{B}$. To reduce the calculation process, we will consider only

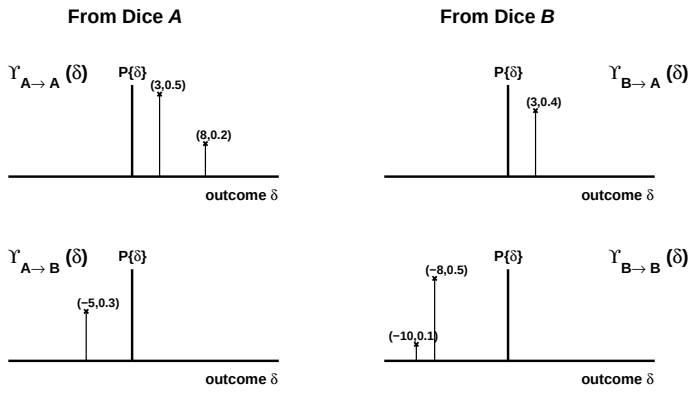

Figure 22: The state transition distributions as prescribed by our dice.

the moments of each of these individual elements about zero. As such we construct a vector, ${}^{n}\underline{x}_{U}$, which takes the form:

$$ {}^{n}\underline{x}_{U} \quad = \quad \left(\begin{array}{c} \sum_{S=-\infty}^{\infty} \varsigma_{1}^{U}(S) S^{n} \\ \sum_{S=-\infty}^{\infty} \varsigma_{2}^{U}(S) S^{n} \end{array} \right) \tag{39} $$

The elements are just the nth moments about zero of the partial distributions within the appropriate state. For $n = 0$ this vector merely represents the probabilities of being in either state at some time $t + U$. We note that for the $n = 0$ case, ${}^{0}\underline{x}_{U+1} = \underline{\underline{T}}\ {}^{0}\underline{x}_{U}$ where $\underline{\underline{T}}$ is the Markov Chain transition matrix as in equation 8. We also note that the transition matrix $\underline{\underline{T}} = {}^{0}\underline{\underline{X}}$ where we define the (static) matrix ${}^{n}\underline{\underline{X}}$ in a similar fashion using the partial distributions (described in figure 22) to be

$$ {}^{n}\underline{\underline{X}} \quad = \quad \left(\begin{array}{cc} \sum_{\delta=-\infty}^{\infty} \Upsilon_{A\to A}(\delta)\delta^{n} & \sum_{\delta=-\infty}^{\infty} \Upsilon_{B\to A}(\delta)\delta^{n} \\ \sum_{\delta=-\infty}^{\infty} \Upsilon_{A\to B}(\delta)\delta^{n} & \sum_{\delta=-\infty}^{\infty} \Upsilon_{B\to B}(\delta)\delta^{n} \end{array} \right) \tag{40} $$

Again, this contains the moments (about zero) of the partial distributions corresponding to the transitions between states. The evolution of ${}^{0}\underline{x}$, the state-wise probabilities with time is trivial as described above. For higher orders, we must consider the effects of superposition and convolution on their values. We know that the for the superposition of two partial distributions, the resulting moments (any order) about zero will be just the sum of the moments of the individual distributions, it is just a summation. The effects of convolution, however must be considered more carefully. The elements of our vector ${}^{1}\underline{x}_{U}$ are the first order moments of the values associated with either of the two states at time $t+U$. The first element of which corresponds to those values of output in state $\boldsymbol{A}$ at time $t + U$. Consider that element

one step later, ${}^{1}x_{1,U+1}$. This can be written as the superposition of the two required convolutions.

$$\begin{aligned}
\sum_{S} \varsigma_1^{U+1}\, S &= \sum_{\delta}\sum_{S} \varsigma_1^{U}\, \Upsilon_{A\to A}\,(S + \delta) + \sum_{\delta}\sum_{S} \varsigma_2^{U}\, \Upsilon_{B\to A}\,(S + \delta) \\
{}^{1}x_{1,U+1} &= {}^{0}x_{1,U}\, {}^{1}X_{1,1} + {}^{0}X_{1,1}\, {}^{1}x_{1,U} + \\
&\qquad {}^{0}x_{2,U}\, {}^{1}X_{1,2} + {}^{0}X_{1,2}\, {}^{1}x_{2,U}
\end{aligned} \tag{41}$$

In simpler terms,

$${}^{1}\underline{x}_{U+1} = {}^{0}\underline{\underline{X}}\, {}^{1}\underline{x}_{U} + {}^{1}\underline{\underline{X}}\, {}^{0}\underline{x}_{U} \tag{42}$$

The $(S + \delta)$ term in the expression relates to the nature of series generating algorithm, $S_{t+1} = S_t + \delta_t$. If the series updating algorithm were altered, this would have to be reflected in this convolution.

The overall output of the system is the superposition of the contributions in each state. As such, the resulting first moment about zero (the mean) for the overall output at U time-steps into the future is simply ${}^{1}\underline{x}_{U+1} \cdot \underline{1}$ where $\underline{1}$ is a vector containing all ones and $\cdot$ is the familiar dot product.

The other moments about zero can be obtained similarly.

$$\begin{aligned}
{}^{0}\underline{x}_{U+1} &= {}^{0}\underline{\underline{X}}\, {}^{0}\underline{x}_{U} \\
{}^{1}\underline{x}_{U+1} &= {}^{0}\underline{\underline{X}}\, {}^{1}\underline{x}_{U} + {}^{1}\underline{\underline{X}}\, {}^{0}\underline{x}_{U} \\
{}^{2}\underline{x}_{U+1} &= {}^{0}\underline{\underline{X}}\, {}^{2}\underline{x}_{U} + 2\, {}^{1}\underline{\underline{X}}\, {}^{1}\underline{x}_{U} + {}^{2}\underline{\underline{X}}\, {}^{0}\underline{x}_{U} \\
{}^{3}\underline{x}_{U+1} &= {}^{0}\underline{\underline{X}}\, {}^{3}\underline{x}_{U} + 3\, {}^{1}\underline{\underline{X}}\, {}^{2}\underline{x}_{U} + 3\, {}^{2}\underline{\underline{X}}\, {}^{1}\underline{x}_{U} + {}^{3}\underline{\underline{X}}\, {}^{0}\underline{x}_{U} \\
\vdots &\quad \vdots
\end{aligned} \tag{43}$$

More generally

$${}^{n}\underline{x}_{U+1} = \sum_{\gamma=0}^{n} {}^{n}C_{\gamma}\, {}^{\gamma}\underline{\underline{X}}\, {}^{n-\gamma}\underline{x}_{U} \tag{44}$$

where ${}^{n}C_{\gamma}$ is the conventional *choose* function.

To calculate time-averaged properties of the system, for example the one-time-step mean or variance, we set the initial vectors such that

$${}^{0}\underline{x}_{0} = {}^{0}\underline{\underline{X}}\, {}^{0}\underline{x}_{0} \tag{45}$$

and ${}^{\beta}\underline{x}_{0} = \underline{0}$ for $\beta > 0$. The moments about zero can then be used to calculate the moments about the mean. The mean of the one time-step increments in output averaged over an infinite run will then be ${}^{1}\underline{x}_{1} \cdot \underline{1}$ and σ^2 will be

$$\sigma^2 = {}^{2}\underline{x}_{1} \cdot \underline{1} - ({}^{1}\underline{x}_{1} \cdot \underline{1})^2 \tag{46}$$

These can be calculated for any size rolling window. The mean of all U-step increments, $S(t+U)-S(t)$) or conversely the mean of the Future-Cast U steps into the future from unknown current state is simply ${}^1\underline{x}_U \cdot \underline{1}$and σ_U^2 will be

$$\sigma_U^2 \;=\; {}^2\underline{x}_U \cdot \underline{1} - ({}^1\underline{x}_U \cdot \underline{1})^2 \tag{47}$$

again with initial vectors calculated from ${}^0\underline{x}_0 = \; {}^0\underline{\underline{X}}\; {}^0\underline{x}_0$ and ${}^\beta\underline{x}_0 = \underline{0}$ for $\beta > 0$. Examining this explicitly for our two dice model, the initial vectors are:

$$\begin{aligned} {}^0\underline{x}_0 &= \begin{pmatrix} \frac{4}{7} \\ \frac{3}{7} \end{pmatrix} \\ {}^1\underline{x}_0 &= \begin{pmatrix} 0 \\ 0 \end{pmatrix} \\ {}^2\underline{x}_0 &= \begin{pmatrix} 0 \\ 0 \end{pmatrix} \end{aligned} \tag{48}$$

and the (static) matrices are:

$$\begin{aligned} {}^0\underline{\underline{X}} &= \begin{pmatrix} 0.7 & 0.4 \\ 0.3 & 0.6 \end{pmatrix} \\ {}^1\underline{\underline{X}} &= \begin{pmatrix} 3.1 & 1.2 \\ -1.5 & -5.0 \end{pmatrix} \\ {}^2\underline{\underline{X}} &= \begin{pmatrix} 17.3 & 3.6 \\ 7.5 & 42. \end{pmatrix} \end{aligned} \tag{49}$$

These are all we require to calculate the means and variances for our system's potential output at any time in the future. To check that all is well, we employ the Future-Cast to generate the possible future distributions of output up till 10 time-steps, Π^1_{char} to Π^{10}_{char}. The means and variances of these are compared to the reduced Future-Cast formalism and also a numerical simulation. This is a single run of the game over 100000 time-steps. The means and variances are then measured over rolling windows of between 1 and 10 time-steps in length. The comparison is shown in figure 23. Fortunately they all concur. The Reduced Future-Cast formalism and the moments about either the mean or zero from the distributions generated by the Future-Cast formalism are identical. Clearly numerical simulations require progressively longer run times to investigate the properties of distributions further into the future, where the total number of possible paths gets large.

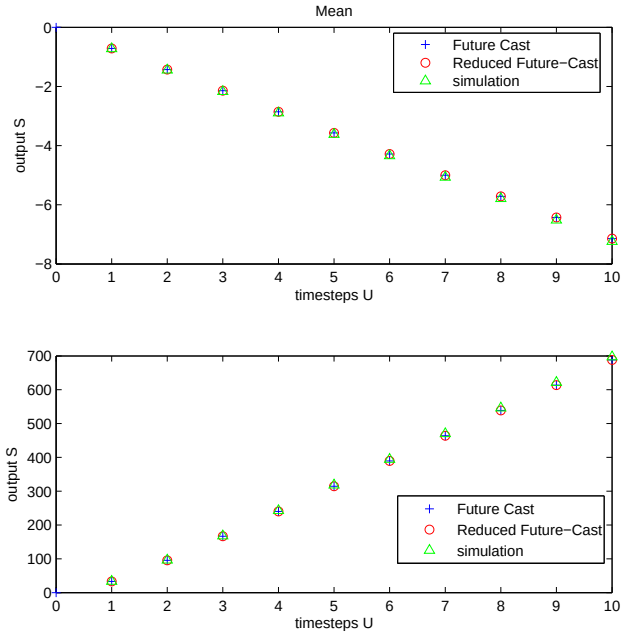

Figure 23: The means and variances of the characteristic distributions Π^{1}_{char} to Π^{10}_{char} as compared to a numerical evaluation and the reduced Future-Cast.

8 Discussion

We have presented an analytical formalism for the calculation of the probabilities of outputs from the B-A-R system at a number of time-steps beyond the present state. The construction of the (static) Future-Cast operator matrix allows the evolution of the systems output, and other macroscopic quantities of the system, to be studied without the need to follow the microscopic details of each agent or species. We have demonstrated the technique to investigate the macroscopic effects of population perturbations but it could also be used to explore the effects of exogeneous noise or even news in the context of financial markets. We have concentrated on single realisations of the quenched disorder matrix, since this is appropriate to the behaviour and design of a particular realization of a system in practice. An example could be a financial market model based on the B-A-R system whose derivatives could be analysed quantitatively using expectation values generated with the Future-Casts. We have also shown that through the normalised eigenvector solution of the Markov Chain transition matrix, we can use the Future-Cast operator matrix to generate a characteristic probability function for a given game over a given time period. The formalism is general

to any time-horizon game and could, for example, be used to analyse systems (games) where a level of communication between the agents is permitted, or even linked systems (i.e. linked games or 'markets'). In the context of linked systems, it will then be interesting to pursue the question as to when adding one 'safe' complex system to another 'safe' complex system, results in an 'unsafe' complex system. Or thinking more optimistically, when can we put together two or more 'unsafe' systems and get a 'safe' one?

We have also presented a simplified and altogether more usable interpretation of the Future-Cast formalism for tracking the evolution of the output variable from a complex system whose internal states can be described as a Markov process. We have illustrated the application of the results for an example case both for the evolution from a known state or when the present state is unknown, to give characteristic information about the output series generated by such a system. The formalism is generalizable to Markov Chains whose state transitions are not limited to just two possibilities and also to systems whose mapping from state transitions to output-space are governed by continuous probability distributions.

Future work will focus on the 'reverse problem' of the broad-brush design of multi-agent systems which behave in some particular desired way – or alternatively, ones which will avoid some particular undesirable behaviour. The effects of any perturbation to the system's heterogeneity could then be pre-engineered in such a system. One possible future application would be to attack the global control problem of discrete actuating controllers[27]. We will also pursue our goal of tailoring multi-agent model systems to replicate the behaviour of a range of real-world systems, with a particular focus on (1) biological and human health systems such as cancer tumours and the immune system, and (2) financial markets.

References

[1] See N. Boccara, *Modeling Complex Systems* (Springer, New York, 2004) and references within, for a thorough discussion.

[2] D.H. Wolpert, K. Wheeler and K. Tumer, Europhys. Lett., **49**(6) (2000).

[3] S. Strogatz, *Nonlinear Dynamics and Chaos* (Addison-Wesley,Reading, 1995).

[4] N. Israeli and N. Goldenfeld, Phys. Rev. Lett. **92**, 074105 (2004).

[5] J.L. Casti, *Would-be Worlds* (Wiley, New York, 1997).

[6] B. Arthur, Amer. Econ. Rev. **84**, 406 (1994); Science **284**, 107(1999).

[7] N.F. Johnson, S. Jarvis, R. Jonson, P. Cheung,Y.R. Kwong, and P.M. Hui, Physica A **258** 230 (1998).

[8] M. Anghel, Z. Toroczkai, K. E. Bassler, G.Kroniss, Phys. Rev. Lett. **92**, 058701 (2004).

[9] S. Gourley, S.C. Choe, N.F. Johnson, and P.M.Hui, Europhys. Lett. **67**, 867 (2004).

[10] S.C. Choe, N.F. Johnson, and P.M. Hui, Phys. Rev. E **70**, 055101 (2004).

[11] T.S. Lo, H.Y. Chan, P. M. Hui, and N. F.Johnson, Phys. Rev. E **70**, 056102 (2004).

[12] D. Challet and Y.C. Zhang, Physica A **246**, 407 (1997);.

[13] D. Challet, M. Marsilli, and G. Ottino, cond-mat/0306445.

[14] N.F. Johnson, S.C. Choe, S. Gourley, T.Jarrett, and P.M. Hui, in *Advances in Solid State Physics* Vol. **44** (Springer, Heidelberg, 2004),p. 427.

[15] We originally introduced the finite time-horizon MG, plus its 'grand canonical' variable-N and variable-L generalizations, to provide a minimal model for financial markets. See M.L. Hart, P. Jefferies and N.F. Johnson, Physica A **311**, 275(2002); M.L. Hart, D. Lamper and N.F. Johnson, Physica A **316**, 649 (2002); D. Lamper, S. D. Howison, and N. F. Johnson, Phys. Rev. Lett. **88**, 017902 (2002); N.F. Johnson, P. Jefferies, and P.M. Hui,*Financial Market Complexity* (Oxford University Press, 2003). See also D. Challet and T. Galla, cond-mat/0404264,which uses this same model.

[16] D. Challet and Y.C. Zhang, Physica A **256**, 514 (1998).

[17] D. Challet, M. Marsili and R. Zecchina, Phys. Rev. Lett. **82**, 2203 (1999).

[18] D. Challet, M. Marsili and R. Zecchina, Phys. Rev. Lett. **85**, 5008 (2000).

[19] See http://www.unifr.ch/econophysics for Minority Game literature.

[20] N.F. Johnson, M. Hart and P.M. Hui, Physica A **269**, 1 (1999).

[21] M. Hart, P. Jefferies, N.F. Johnson and P.M. Hui, Physica A **298**, 537 (2001).

[22] N.F. Johnson, P.M. Hui, Dafang Zheng, and M. Hart, J. Phys. A: Math. Gen. **32**, L427 (1999).

[23] M.L. Hart, P. Jefferies, N.F. Johnson and P.M. Hui, Phys. Rev. E **63**, 017102 (2001).

[24] P. Jefferies, M. Hart, N.F. Johnson, and P.M. Hui, J. Phys. A: Math. Gen. **33**, L409 (2000).

[25] P. Jefferies, M.L. Hart and N.F. Johnson, Phys. Rev. E **65**, 016105 (2002).

[26] N.F Johnson, D.M.D. Smith and P.M. Hui, Europhys. Lett. (2006) in press.

[27] S. Bieniawski and I.M. Kroo, AIAA Paper 2003-1941 (2003).

[28] I.M. Sobol, *A Primer for the Monte Carlo Method* (Boca Raton, CRC Press, 1994).